Overcoming the Memory Wall

Center for Applied Scientific Computing

Overview

CPU speeds double approximately every eighteen months, while main memory speeds double only about every ten years. These diverging rates imply an impending "Memory Wall," in which memory accesses dominate code performance. The Center for Applied Scientific Computing (CASC) at Lawrence Livermore National Laboratory (LLNL), in collaboration with the University of Utah, is extending dynamic access optimizations (DAO) mechanisms to symmetric multiprocessors (SMPs). DAO techniques improve memory performance by changing the order or apparent location of memory accesses. Implementing these mechanisms in commercial systems could effectively triple LLNL's computer purchasing power.

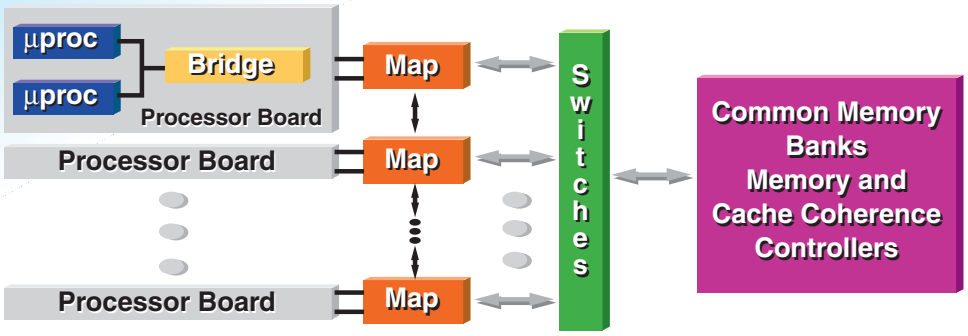

Figure 1. The FPGA memory controllers and MAPS of the SRC-6 are an ideal test bed for SMP-based DAO.

Standard memory access techniques do not use memory bandwidth effectively. Many researchers anticipate a Memory Wall in which memory accesses imply an absolute performance limit, and improvements in CPU speed provide no performance benefit. Dynamic access optimization (DAO) mechanisms can use the memory system more effectively by changing the order or apparent location of memory accesses from those generated by the issuing program to ones that use the memory system more effectively without changing the results, thus lowering the Memory Wall.

This project is extending DAO mechanisms to consider additional memory access issues, such as cache coherency, that arise in symmetric multiprocessors (SMPs). The deliverables of this project include recommended improvements to the compilers, operating systems, and firmware of SMP-based systems. These modifications could be of significant benefit to scientific codes and reduce the run times of many LLNL codes by a factor of three or more without any source code modifications.

DAO techniques have shown significant promise to overcome the Memory Wall without requiring complex source code changes. These techniques use existing memory hardware more effectively. For example, altering the execution order can exploit memory hardware characteristics such as interleaved memory banks and hot dynamic random access memory (DRAM) rows. More advanced techniques that alter the apparent location of accesses reduce cache miss rates as well as lowering memory bandwidth demand and reducing an application's cache footprint. DAO mechanisms can reduce run times of memory intensive portions of programs by factors of 2 to 13. Previous projects investigating DAO focus on uniprocessor systems and require special-purpose hardware.

Dynamic Access Optimizations in Uniprocessors

Results for DAO techniques in uniprocessor systems are promising. The Stream Memory Controller (SMC) combines compile-time detection of streams (i.e., predictable access patterns such as strided array accesses) with hardware that modifies the execution order of the memory accesses to exploit memory hardware characteristics. Prototypes of the special purpose hardware were implemented and interfaced to an Intel i860 processor. Results for benchmark kernels show that the SMC consistently improves memory latency and bandwidth utilization for regular computations by a significant factor. In fact, it increases the effective memory bandwidth delivered to some computational kernels by an order of magnitude.

The Impulse memory controller also employs DRAM scheduling techniques like those used in the SMC. However, it also uses shadow memory, which is a more advanced technique that can optimize both cache and main memory use. Shadow memory can be used to target irregular access patterns, such as array accesses that use an indirection vector. Specialized hardware and firmware gather groups of indirect accesses into a cache line in order to increase spatial locality. Thus, shadow memory allows a single main memory access to accomplish what would normally take several main memory accesses. As a result, shadow memory is one of the few approaches to improve memory system performance for irregular applications, such as sparse matrix-vector products. Implementing shadow memory requires minor changes to the virtual memory module of the operating system, while the compiler and other run time

UCRL-TB-145881

URL: http://www.llnl.gov/CASC/MemWall

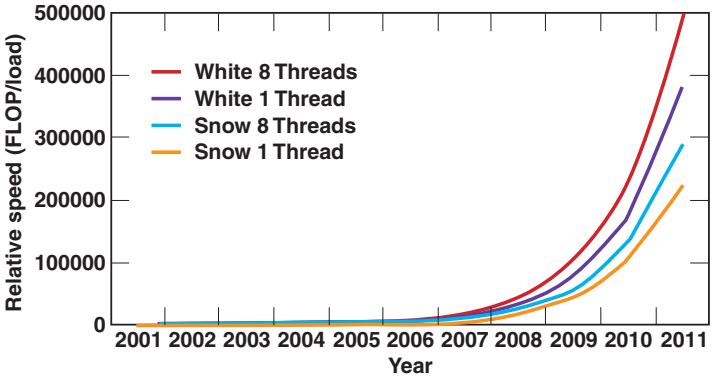

Figure 2. Processor and memory speeds are diverging based on memory sysems performance in current SMPs, as many as one-half million floating point operations could be completed during one main memory access on projected SMPs of 2011.

requirements are similar to those of the SMC. Uniprocessor simulations of the Impulse memory controller demonstrate improvements in run time of up to a factor of 3.47.

Dynamic Access Optimizations in Symmetric Multiprocessors

Although promising, DAO techniques for uniprocessors do not target the systems in use at LLNL. All major LLNL computing resources are clusters with SMP nodes. Thus, we need DAO techniques that support simultaneous access to the memory system by multiple processors. Difficulties arise in SMPs for both types of DAO techniques, access reordering and shadow memory. This project will address these difficulties.

Applying access reordering to SMP-based systems could result in the accesses of some processors being delayed in order to improve performance for one processor, which would decrease overall performance. Thus, SMP-aware DAO mechanisms must guarantee fairness. However, simple solutions may work well because the interleaved accesses of the multiple processors can fill the physical memory space more sufficiently. Therefore, accesses can use the memory system hardware even more efficiently than in uniprocessor systems.

Relationships to Emerging Computer Architecture Trends

Current architecture trends recognize the increasing importance of memory system performance. We have established a close relationship between DAO techniques and reconfigurable computing, a significant emerging trend in computer architecture. As part of this project, LLNL has established a close working relationship with SRC Computers, an industry leader in reconfigurable computing. We expect our exploration of DAO techniques for multi-adaptive processors (MAPs) to result in an important mechanism for SRC systems.

An important trend in computer architecture is to locate multiple processor cores on the same chip. The presence of multiple cores on the same chip will increase the rate of memory accesses, thus significantly increasing memory bandwidth requirements. As these SMP chips become prevalent, SMP-aware DAO mechanisms will prove essential to reasonable system performance.

Three-Year Plan

This project, in collaboration with Sally McKee of the University of Utah and Frank Mueller of North Carolina State University, is extending DAO to SMPs. The goals of this research effort are to:

- Extend DAO mechanisms to SMP-based systems
- Implement DAO in a commercially available system
- Demonstrate the benefit of SMP-aware DAO mechanisms for current LLNL systems
- Demonstrate the relevance of SMP-aware DAO techiques to future systems
- Significantly improve the run-time and CPU utilization of LLNL large memory codes.

These goals require a comprehensive three-year research plan. During the first year of the project, we are establishing the infrastructure required for our research. In the second year, we will simulate and implement the first SMP-aware DAO techniques and refine the techniques based on our initial results. We will investigate the performance of SMP-aware DAO techniques in the presence of message passing memory traffic and explore techniques that target systems-on-a-chip (SOC) and processor-in-memory (PIM) in the third year of the project, as well as continuing the iterative process of the second year.

We are collaborating with Sally McKee of the University of Utah and Frank Mueller of North Carolina State University. For more information about this work, contact: Bronis de Supinski, (925)422-1062, bronis@llnl.gov.

This work was performed under the auspices of the U.S. Department of Energy by University of California Lawrence Livermore National Laboratory under contract No. W-7405-Eng-48.